Australian Government
Department of Defence
Defence Science and
Technology Organisation

Cumulative Track-Initiation Probability as a Basis for Assessing Sonar-System Performance in Anti-Submarine Warfare

M.P. Fewell and J.M. Thredgold

Maritime Operations Division
Defence Science and Technology Organisation

DSTO-TN-0932

ABSTRACT

Calculations of sonar-system performance for anti-submarine warfare yield detection probability as a function of range to the target; whereas sonar operators typically think in terms of a ‘detection range’. This Note considers how to connect these two concepts for active sonar. Four types of probability are explored and their relative advantages teased out in the light of real-world examples. It is concluded that cumulative probability of track initiation provides the most practical route to a definition of detection range, for two reasons. First, unlike detection probability as usually computed, it produces values that operators are likely to regard as believable. Secondly, it reflects more closely than the standard detection probability the steps that operators typically go through in deciding to declare a detection.

RELEASE LIMITATION

Approved for public release

Published by

Maritime Operations Division
DSTO Defence Science and Technology Organisation
PO Box 1500
Edinburgh South Australia 5111 Australia

Telephone: (08) 8259 5555
Fax: (08) 8259 6567

AR-014-676
December 2009

Cumulative Track-Initiation Probability as a Basis for Assessing Sonar-System Performance in Anti-Submarine Warfare

Executive Summary

The key quantity that an operator wishes to know about any surveillance asset is its detection range—how far away a target of interest can be detected. Sonar systems for anti-submarine warfare are no exception; performance is often specified in terms of an achievable detection range. However the concept of a fixed 'range of the day' does not sit easily with the realities of highly variable acoustic propagation and ambient noise in the sea. Real-world variability and fluctuations are more readily expressed in terms of detection probability as a function of target range and depth. This report examines ways of deriving detection range from detection probability for active sonar through a consideration of four types of probability:

- single-ping detection probability p_d, which is the result of acoustic-propagation calculations,
- cumulative detection probability P_d, which is the probability of obtaining a detection on the current ping or from any previous ping since the start of the scenario vignette,
- 'local' track-initiation probability p_{ti}, which is the probability of starting a track after 5 consecutive pings, assuming that the track-initiation rule requires 3 or more detections in the 5 pings, and
- cumulative track-initiation probability P_{ti}, which is the probability that a track is started on the current ping, or was started on any previous ping.

Our analysis concludes that P_{ti} provides the best basis for defining detection range. The advantages and disadvantages of P_{ti} compared with the other probabilities are summarised in the traffic-light diagram on the next page. In detail:

- Cumulative probabilities, whether P_{ti} or P_d, require a scenario for their calculation. Hence their values are more scenario dependent than p_d or p_{ti}.
- Cumulative probabilities provide a basis for comparing sonar systems with different ping rates or transmission modes, unlike p_d or p_{ti}.
- Single-ping p_d, by definition, does not allow the analysis of the effect of a long series of pings, unlike the cumulative probabilities. Local p_{ti} considers groups of 5 pings only (or n pings where the track-initiation rule is m-in-n).
- Track-initiation probabilities provide some recognition of the problem of false detections, unlike detection probabilities, while skirting around the difficult issue of data association.
- As regards several sonar systems acting together and pooling data, p_d provides no basis for analysis. Combinations of several p_d values are required, as is the

situation the other three probabilities. However, P_d permits the analysis of track-level data fusion only, whereas p_{ti} and P_{ti} also allow the analysis of detection-level fusion.

- Since cumulative probabilities never decrease with time, they provide an unambiguous basis for defining detection range, unlike p_d or p_{ti}.
- The actual definition of detection range, however, requires the choice of a probability value, which is unavoidably arbitrary. All probabilities have this deficiency, but it is worst for cumulative detection probability P_d, because the 'natural' choice of 50% gives unrealistically large detection ranges in most environments. On the other hand, experience suggests that P_{ti} = 50% works well, at least for persistent sonar systems like ship-borne hull-mounted sonars and towed arrays.

The purpose of any piece of military operations analysis is to inform a decision maker. The success of the analysis is closely related to how useful the client finds it. Analytical techniques based on cumulative track-initiation probability have been used to inform decisions in support of a major defence acquisition project, to the satisfaction of the client.

Summary of strengths and weaknesses of various probabilities for the purposes of defining the 'detection range' of a sonar system.

Probability:	single-ping detection probability	local track-initiation probability	cumulative detection probability	cumulative track-initiation probability
independent of scenario*	🟢	🟡	🔴	🔴
allows comparison of sonars with different ping rates or modes	🔴	🟡	🟢	🟢
provides some recognition of the influence of false detections	🔴	🟢	🔴	🟢
includes effect of many repeated pings	🔴	🟡 5 pings only	🟢	🟢
provides a basis for analysing networks of sonar systems	🔴	🟢	🟡 track-level fusion only	🟢
unambiguous definition of detection range	🔴	🔴	🟢	🟢
no arbitrariness in definition of detection range	🟡 p_d = 50% acceptable	🟡 p_{ti} = 50% acceptable	🔴 P_d = 50% unrealistic	🟡 P_{ti} = 50% acceptable

*This refers to the disposition and motion of the platforms in the scenario. All probabilities are absolutely dependent on the environment, time of year, nature of the target and target depth.

Authors

M.P. Fewell
Maritime Operations Division

Matthew Fewell joined DSTO in 2001, coming from an academic physics background. He has worked and published in experimental nuclear structure physics, gaseous electronics, atom–photon interactions including coherent effects, laser physics, plasma processing of materials, the theory of network-centric warfare and its modelling (including cognitive issues), human-in-the-loop experimentation, and weapon–target allocation in ship air defence. He is at present working on the analysis of anti-submarine warfare from surface-ship and maritime-patrol perspectives.

J.M. Thredgold
Maritime Operations Division

Jane Thredgold commenced working in the Maritime Operations Division at DSTO in early 2007, after studying for a PhD in Mathematics at the University of South Australia. She is currently performing operations analysis work relevant to anti-submarine warfare.

Contents

Figures

Tables

Acronyms

ADF	Australian Defence Force
ASW	anti-submarine warfare
AWD	air-warfare destroyer

1. Introduction

The key quantity that an operator wishes to know about any surveillance asset is its detection range—how far away a target of interest can be detected. Sonar systems for anti-submarine warfare (ASW) are no exception; performance is often specified in terms of an achievable detection range. Unfortunately this leads to difficulties in the case of sonar; for the concept of a fixed ‘range of the day’ does not sit easily with the realities of highly variable acoustic propagation and ambient noise in the sea [1], for reasons that are illustrated in §2.

This note advocates the use of ‘cumulative probability of track initiation’ P_{ti} as a route to the computation of realistic and operator-believable active sonar detection ranges from the results of acoustic-propagation calculations. The concept is not new—P_{ti} has been used for ASW analysis at least once before [2]—but the advantages of its use are not as widely recognised as they deserve to be.

Cumulative track-initiation probability is defined in §2.2.1, where an algorithm for its calculation is described and its advantages are canvassed. Techniques similar to those described in §2.2.1 were used to inform decisions in support of a major defence acquisition, to the satisfaction of the client.

2. Connecting Detection Probability and Detection Range

2.1 Single-Ping Detection Probability

Calculations of the performance of ASW sonar systems typically give results as single-ping detection probability p_{d}. This is the probability that the system, after pinging once, will register an echo from the submarine as a detection. State-of-the-art sonar-system performance calculations are very complicated, taking into account many details of underwater acoustic propagation, the ambient noise field in the sea and the operation of the acoustic processing system. Figure 1(a) shows a typical result: p_{d} is colour-coded as a function of the target submarine’s range and depth.[a] Figure 1(b) shows part of the data in (a) in a more conventional way, as a plot of p_{d} versus range for a particular submarine depth, in this case 50 m.

---

[a] Detection probability may also be a function of target bearing and aspect angle; these refinements are omitted here, but their inclusion causes no insurmountable complication for the main point of this Note, the calculation of P_{ti} .

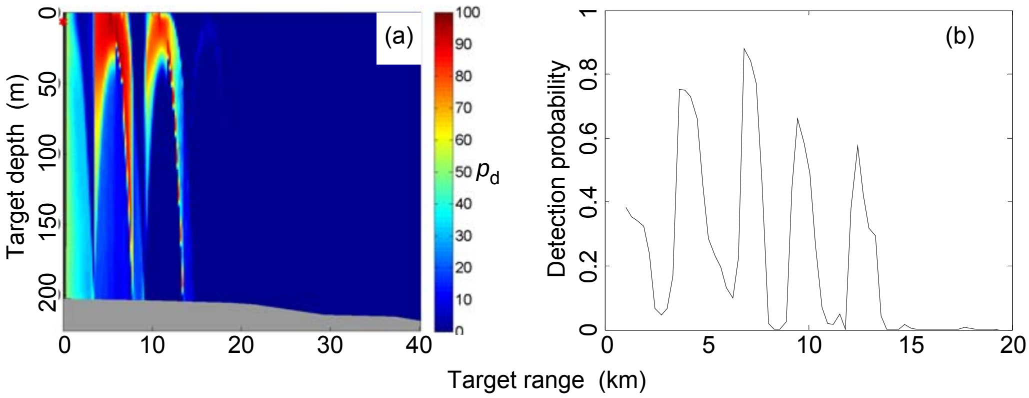

Figure 1: Example of results of sonar-system performance calculations for a notional realistic sonar system, geographic location, time of year and target: (a) colour-coded contour plot of single-ping detection probability p_d *as a function of target range and depth and (b) the same as in (a), but a conventional plot of* p_d *against target range for a particular target depth (50 m in this case). The grey region at the bottom of panel (a) shows the profile of the sea floor. Note the very different length scales on the two axes in this panel.*

It should be carefully noted that the term 'detection' is used here in an engineering sense: it means no more than that the target of interest has returned sufficient acoustic energy to cause the sonar system to display a spot on a screen. This is often rather less than what operators mean by 'detection', which includes elements of identification and the elimination of false sonar returns.

Sometimes, a detection range is derived from a plot like Figure 1(b) by picking a p_d value, say 50%, and adopting the range at which p_d has that value. Several criticisms can be levelled at this method:

1. The p_d value chosen is arbitrary; why pick 50%?
2. Because p_d can increase as well as decrease with range, there may be more than one range at which its value is 50%. Figure 1 was chosen as a severe example of this: it shows a case of shallow water with the acoustic energy bouncing between bottom and surface several times. Which of the range values at which p_d equals 50% should be chosen as the detection range?
3. The quantity p_d is a single-ping value, whereas almost always sonar systems ping repeatedly. This operational reality should be taken into account, not only for its own sake, but also because, when comparing different sonar systems, using a metric based on p_d may result in biased decisions if the different systems have different ping rates or patterns of transmission (e.g. omni-directional compared with sector compared with ripple-fire[b] etc.)
4. It is rarely the case that a single detection, in the engineering sense of the term, is sufficient for an operator to declare a detection, because of false alarms. Although a specified false-alarm rate is part of the input to calculations like those shown in Figure 1, the use of p_d alone, without explicit reference to false-alarm rate, ignores the operational

[b] 'Ripple fire' refers to a sonar mode in which transmissions are sent out sequentially in two or more directions prior to listening. The rationale is that this puts more energy into the water in a given direction than would be possible with a single transmission spread simultaneously over all the directions of interest.

impact of false detections that are difficult to distinguish from real detections and that are typically produced at the rate of several per ping.

5. If more than one sonar system is available and operating, then ASW effectiveness can be increased through appropriate integration of detection information [3]. This is difficult to assess on the basis of single-ping p_{d} alone.

Points 1–3 have long been acknowledged; points 4 and 5, though probably understood, have rarely been taken into account in past work. Points 2 and 3 are standardly addressed by moving from single-ping p_{d} values to a cumulative detection probability P_{d}, as described in the next subsection.

2.2 Cumulative Detection Probability

The cumulative detection probability P_{d} is the probability of obtaining a target detection after a series of pings.(c) Its calculation involves combining several p_{d} values, so a scheme for selecting the actual values to be combined is required. In other words, a scenario must be specified. To illustrate how this works, a very simple scenario is introduced in §2.2.1. Subsections 2.2.2 and 2.2.3 detail its use to compute P_{d} values and, from there, detection range.

2.2.1 Example Scenario

Figure 2 shows schematically a simple scenario involving a surface ship transiting at constant heading and speed. A threat submarine ahead of the ship detects it at bearing b (as measured from the ship) and range R, and then tracks toward intercept. The submarine is assumed to transit at half the speed of the surface ship and to set a course so that bearing b remains constant as the two vessels close. A surface-ship speed v_{s} of 15 kn was chosen, and therefore the submarine speed v_{u} is 7.5 kn. The bearing b was chosen to be 20°.(d)

For simplicity, it is assumed that the surface ship pings regularly during the scenario, using just one sonar mode.(e) The scenario should start at a range R_{initial} where the single-ping p_{d} values are negligibly small.(f) That is, p_{d} values as a function of R are required as inputs to the calculation. The scenario can be stopped at any convenient value of R; in the examples presented below, we stopped at R = 1 km.

(c) In this note, we consistently use lower case p to denote single-ping, or few-ping, or 'instantaneous' or 'local' probabilities and upper case P to denote cumulative probabilities. (We did not follow this convention in earlier publications, e.g. [3].)

(d) The maximum possible bearing for the speeds chosen is 30°. That is, for b > 30°, intercept is possible only if the submarine increases its speed.

(e) More complicated search schemes — for example, sector sweeps at a regular sequence of bearings, or some initial short pings to clear the blind zone — can be readily incorporated at no more cost than the additional book-keeping required.

(f) Uncritical adherence to this requirement can lead to a difficulty with excessively long ping intervals, as discussed in the next subsection.

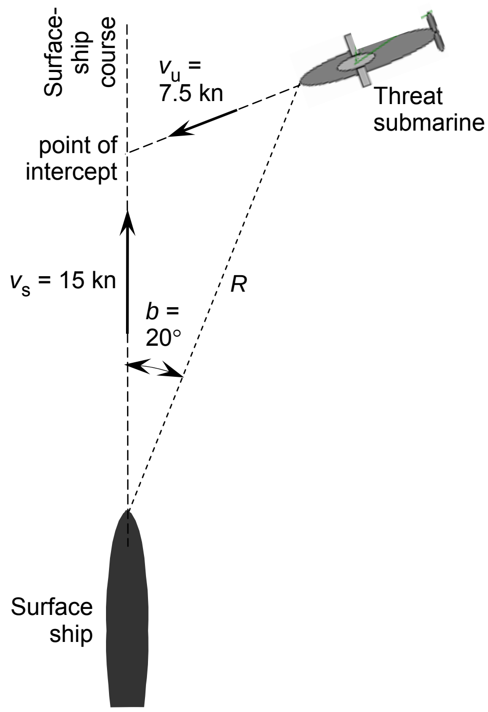

Figure 2: Schematic diagram of the scenario used in the calculation of cumulative detection probabilities

2.2.2 Computation of P_d

The cumulative detection probability $P_{d,K}$ after K pings is given by

$$P_{d,K} = 1 - \prod_{k=1}^{K} \left(1 - p_{d,k}\right), \tag{1}$$

where $p_{d,k}$ is the single-ping detection probability for ping k. In practice, it is efficient to compute $P_{d,K}$ using the equivalent recurrence relation:

$$P_{d,K} = 1 - \left(1 - P_{d,K-1}\right)\left(1 - p_{d,K}\right) \quad \text{(for } K > 1\text{).} \tag{2}$$

With the initial value of $P_{d,1} = p_{d,1}$. Ping 1 is the first ping following scenario start at range R_{initial}. At each subsequent ping thereafter, P_d accumulates according to the recurrence relation. Results for two p_d versus range curves are shown in Figure 3.

The range at which to start the scenario can be a significant modelling issue. Some environments show significant levels of instantaneous detection probability at large distances, for example, when a convergence zone is present. However, the low speed of sound in water means that long look distances require correspondingly long inter-ping intervals. For example, the out-and-back acoustic propagation time is 27 s over 20 km, 40 s over 30 km and 108 s over 80 km, using a speed of sound in water of 1500 m s^{-1}. If the scenario is started at 80 km range, then the maximum ping rate is one quarter of that if the scenario is started at

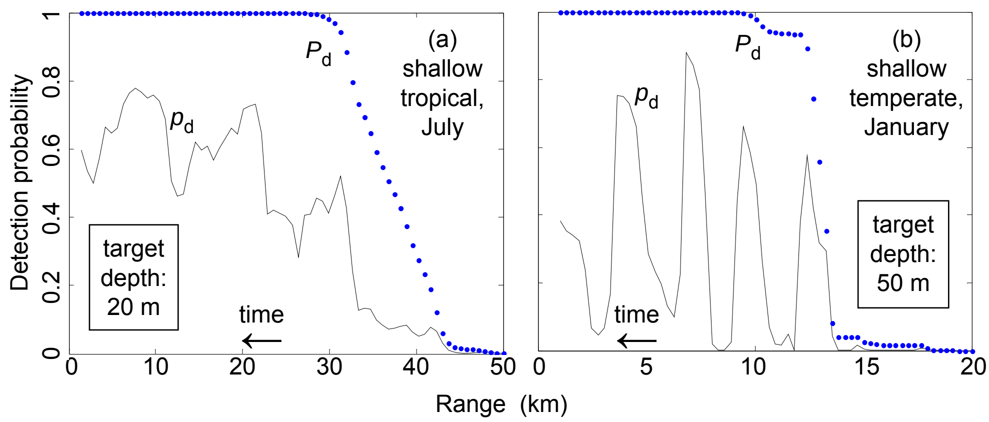

Figure 3: *Values of single-ping detection probability* p_d *(full black line) and cumulative detection probability* P_d *(dotted blue line) for two environments and target depths. Accumulation of* P_d *began at ranges of (a) 50 km and (b) 20 km. The range value at each blue dot shows the range applying at the corresponding ping.*

20 km range. This can cause P_d to accumulate slowly.[g] Clearly, there is a trade-off between ping rate and utilising whatever instantaneous p_d may be available at large range [4]. This issue should be addressed on a case-by-case basis in each environment where it arises, having regard, among other things, to likely operational practice with the sonar system being modelled.[h]

2.2.3 Detection Range from P_d

If it is desired to quote a detection range, then once again one must select a probability value for the purpose. As Figure 3 illustrates, cumulative P_d often rises rapidly, reaching values close to 1.0 at quite large ranges. For this reason, 50% P_d can seem too low a value for the definition of detection range. In the cases shown in Figure 3, one would be tempted to choose a P_d value over 90% so that the resulting detection range may be judged to be 'realistic'.

2.2.4 Critique and Issues

Cumulative probabilities are monotonic; that is, they do not decrease with the passage of time. This eliminates issue 2 on p. 2: choosing a P_d value defines a unique detection range. Issue 3 is also dealt with in principle, since any characteristic of a given sonar system — such as ripple-fire modes (footnote b, p. 2), restrictions on ping rates, the use of sector searches or short pulses to clear a blind zone, the effect of employment tactics with dipping sonars, etc. — can be included as part of the scenario, so affecting the rate at which P_d accumulates.

Issue 5 (p. 3) can be handled in principle too, by extending the definition of cumulative detection probability. If detections from two sonar systems are being combined then,

[g] In principle, this limitation could be overcome by coding pings. However, it would also be necessary to stagger inter-ping spacings in order to search the ring-shaped blanking zone arising from the fact that monostatic sonars cannot listen while transmitting. Fielded monostatic sonar systems typically do not have the capability of interleaving pings.

[h] For example, if the environment has a convergence zone, would operators seek to exploit it?

assuming that the sonars are acting independently, the overall cumulative detection probability of the suite of sonar systems is

$$P_{\mathrm{d,sonar\,suite}} = 1 - \left(1 - P_{\mathrm{d,sonar\,1}}\right)\left(1 - P_{\mathrm{d,sonar\,2}}\right), \tag{3}$$

with an obvious generalisation if there are more than two sonar systems contributing to the detector suite.

On the other hand, issue 1 on p. 2 remains: one must still choose a P_{d} value at which to declare a 'detection range', and any such choice is necessarily arbitrary. If anything, the issue is more acute with P_{d} than p_{d} because the 'natural' choice of 50% almost always gives unrealistically large detection ranges when applied to P_{d}.

Finally, issue 4 on p. 2 has been ignored. Typical false-alarm rates from sonar systems mean that the detection of a distant submarine might not be recognised among the clutter. It is usually necessary to detect the submarine more than once in order to distinguish real from false detections. This consideration exposes the problem with $P_{\mathrm{d},K}$, which gives the cumulative probability of just a single detection from any of the K pings. This is why $P_{\mathrm{d},K}$ rises so rapidly, and is also the reason why automated tracking algorithms typically require, say, 3 detections in 5 consecutive pings before a track is initiated.

The way forward is now clear: we propose carrying the analysis a little further along the ASW kill chain to compute probabilities connected with track initiation, as described in the next section.

3. Probability of Track Initiation

The discussion of §2 shows the advantages of using a cumulative probability. The most conceptually straightforward way of constructing this is through an 'instantaneous' or 'local' track-initiation probability. This is described in §3.1. The cumulative track-initiation probability could in principle be computed from this, but it turns out to be much easier to calculate it directly from detection performance using a Monte-Carlo method, as described in §3.2.2. Subsection 3.2.3 gives some examples comparing the behaviour of the two cumulative probabilities—of detection and of track initiation.

3.1 'Local' Track-Initiation Probability

Automated tracking systems typically apply a rule such as: 'start tracking an object if it is detected m or more times in n opportunities'. For a given value of m, the 'local' track-initiation probability p_{ti} is:

the probability that a track is initiated after exactly n pings.

This is sometimes called an 'instantaneous' probability, but the term is misleading because the probability refers to the period of time over which the n pings are emitted, which is typically many minutes. Although this length of time is usually small compared to the

duration of a scenario, we prefer to describe the probability as 'local', which we mean in a temporal sense.

Use of p_{ti} is common in studies of radar (e.g. [5] and refs. therein). For sonar, common values of *m* and *n* are 3-in-5, which are adopted here for definiteness.(i) If all five pings have the same detection probability p_{d} and are statistically independent, then a simple expression for p_{ti} results [3]:

$$p_{\mathrm{ti}} = p_{\mathrm{d}}^{3}\left(10 - 15p_{\mathrm{d}} + 6p_{\mathrm{d}}^{2}\right). \tag{4}$$

This expression forms the basis for a coverage-area analysis [3], but it is not useful in the present context for two reasons. First, it is most unlikely in any realistic scenario that p_{d} would be the same for all five pings. It is possible, using the rules for combining probabilities, to write an expression like Equation (4) for a situation where the $p_{\mathrm{d},k}$ values ($1 \le k \le 5$) are all different, but the result is too complicated to be useful and much too complicated to generalise to other values of *m* and *n*.

The second problem with Equation (4) concerns a significant aspect of data processing in real sonar systems, the 'data-association problem'. It is assumed in deriving Equation (4) that every detection is known with certainty to be a detection of the target of interest. Real-world ASW sonar systems, on the other hand, typically return several false detections with every ping. This is because operators almost always need to detect submarines at ranges approaching the maximum possible, which requires them to set the detection threshold to a low value. Many false detections arising from random noise spikes exceeding the threshold is the inevitable consequence. Since a false detection can look just the same as a real detection, which of the detections to associate with the target can be a significant problem. We are currently exploring the effect of false detections on tracking, including track initiation [6,7]. For present purposes, the data-association problem can be resolved in one of two ways:

- Perhaps false detections are not, in fact, identical to real detections. They might display some characteristic that allows an operator to distinguish false from real, such as, for example, the lack of Doppler shift or a distinctive tonal quality in the returning echo.
- If not, then the problem can be circumvented, though only formally, by re-interpreting p_{d} to mean the probability of detection *and* correct data association (as with, for example, the use of the product $P_{\mathrm{D}}\,P_{\mathrm{G}}$ in Reference 8). The difficulty here is that acoustic-propagation calculations, which provide values of detection probability like those shown in Figure 1(a), do not provide estimates of the data-association probability.

Another complication arises from the fact that usually many more than *n* pings are emitted during a scenario.(j) We deal with this by adopting the sliding-window approach illustrated in Figure 4. Groups of 5 consecutive pings are considered. Each group has a p_{ti} value depending on the $p_{\mathrm{d},k}$ values for the 5 pings concerned. With this further consideration, there would

(i) Examples of the effect of choosing other values are given in ref. [3].

(j) Exceptions to this may occur with sonobuoys or dipping sonars. For example, a dipper may be pinged only a few times before relocation. Examples of this sort should be treated case by case on their merits.

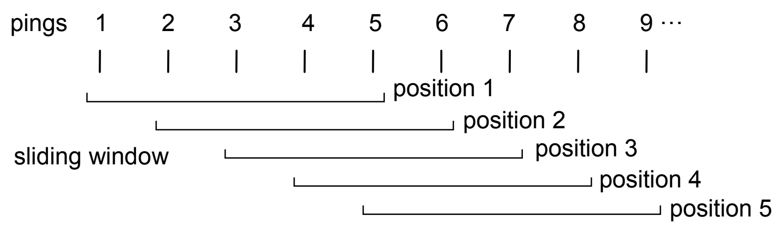

Figure 4: Sliding-window concept for defining a 'local' track-initiation probability using the 3-in-5 rule when more than 5 pings have been emitted. The vertical lines represent the pings and the long brackets represent successive applications of a sliding window. The track initiation rule is defined to be satisfied with the first window that contains 3 detections.

appear to be little point in pursuing the goal of writing down a general algebraic expression for $p_{\mathrm{ti},j}$, the track-initiation probability for the sliding window in position j. Another approach must be sought. Our suggestion is presented in §3.2.2, following the definition of cumulative track-initiation probability.

3.2 Cumulative Track-Initiation Probability

3.2.1 Definition

The cumulative track-initiation probability P_{ti} is:

the probability that a track is initiated on the current ping, or has already been initiated on a previous ping.

If expressions were available for $p_{\mathrm{ti},j}$, the local track-initiation probability for the sliding window in position j, then these could be used to compute $P_{\mathrm{ti},J}$, the cumulative track-initiation probability after J positions of the sliding window had been considered, using an analogous expression to Equation (1). However, values of $p_{\mathrm{ti},j}$ are not available for reasons discussed in §3.1,[k] so we proceed as described in the next subsection.

3.2.2 Monte-Carlo Calculation

Cumulative P_{ti} can be calculated by a Monte-Carlo method. A scenario such as that described in §2.2.1 is run, say, 1000 times. Each run has exactly the same number of pings and any given ping in a run occurs at exactly the same range as corresponding pings in other runs. At each ping in each run, a random number is drawn from a uniform distribution on the interval [0, 1] and compared with the local p_{d} value for the appropriate range to determine whether that ping resulted in a detection. In this manner, each run produces a list of pings marked as either detections or misses. Figure 5 shows schematically part of the set of runs for a particular sonar system, environment, target type and target depth, with ticks marking detections and crosses misses.

[k] Castella [9] gives recurrence relations for both $p_{\mathrm{ti},j}$ and P_{ti} for all cases with $n \leq 4$ but, not only are these difficult to generalise to larger values of n, also he assumes equal p_{d} for all pings. Williams [10] considers varying p_{d}, but some of his expressions appear to be erroneous (in particular his Equations 2 and 3).

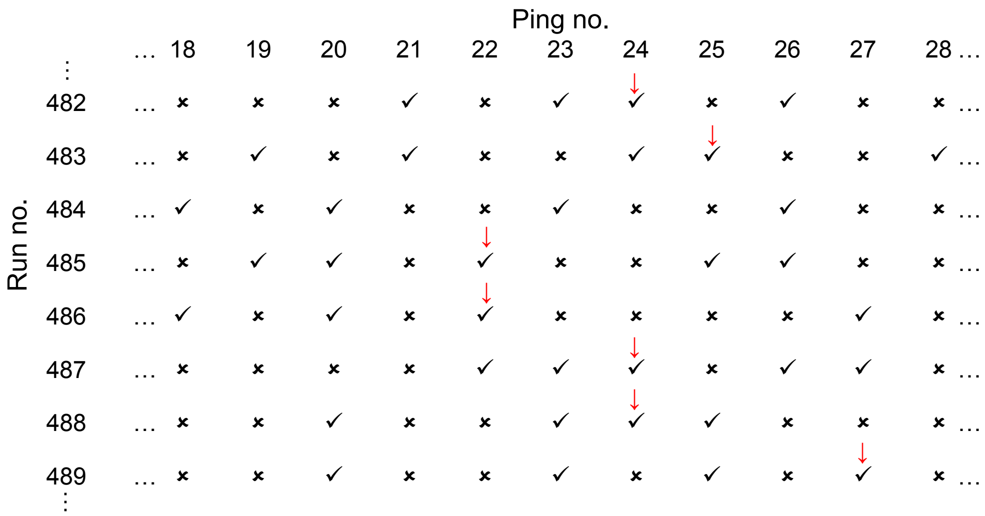

Figure 5: *Schematic representation of the result of a Monte-Carlo calculation for a given sonar, environment, target type and target depth. Parts of 8 of the 1000 runs are shown. Each run has exactly the same number of pings, with detections being determined randomly according to the instantaneous p_d value applying at each ping. Ticks represent detections and crosses misses. Red arrows show pings at which the 3-in-5 track initiation rule is first satisfied, assuming that it was not satisfied prior to ping 18. Run 484 has not yet achieved track initiation.*

Once the 1000 runs are computed, they are scanned for track initiation. That is, the sliding window is applied and the ping number at which track initiation is first achieved is marked as shown schematically by the red arrows in Figure 5.

To determine $P_{\mathrm{ti},J}$ after a given ping J, the number of runs that have achieved track initiation on or before that ping is determined. That is, one counts red arrows down column J in Figure 5 and adds to this the total number of red arrows in all columns to the left of column J. This number expressed as a fraction of 1000 is an estimate of the cumulative P_{ti} following ping J. The accuracy of the estimate is related to the number of runs and can be explored by computing several sets of 1000 runs.

3.2.3 Examples of P_d and P_ti

Figure 6 shows four examples of values of p_d (black lines),[(1)] P_d (blue dots) and P_ti (red crosses). Each dot or cross corresponds to a ping in the scenario. The accumulation begins at the scenario start range, which in each case is the maximum range shown on the abscissa. The scenario involves closing platforms, so range becomes smaller as time progresses. As can be seen, P_d reaches quite high values very early in each scenario, at relatively large ranges.

The rapid increase in P_d as the scenario progresses arises, of course, because the submarine need be detected just once to fulfil the success criterion for P_d. Requiring instead 3 detections in 5 consecutive pings overcomes this problem, as the examples in Figure 6 show.

[(1)] The p_d values shown do not take account of the blind zone near zero range, the extent of which depends on the pulse length transmitted. This may affect the accumulation of P_ti near the end of the scenario; in particular, the rise in P_ti values apparent at short range in panels (c) and (d) will not occur when long pulses are used.

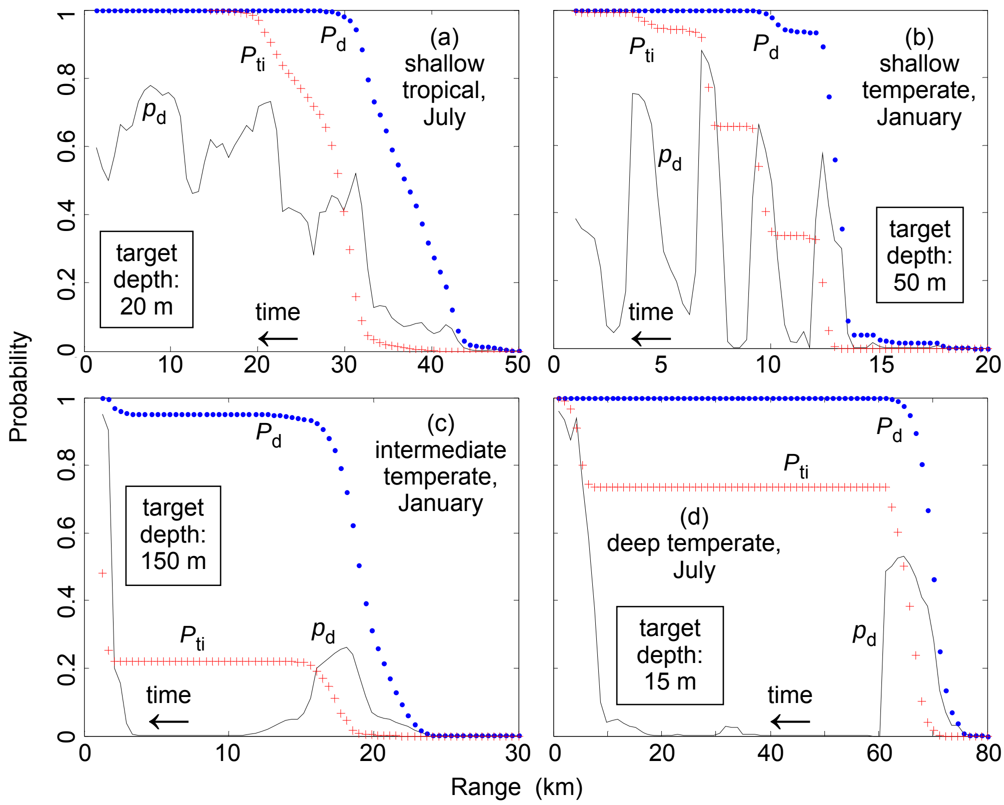

Figure 6: *Curves of instantaneous detection probability* p_d *(black lines), cumulative detection probability* P_d *(blue dots) and cumulative track-initiation probability* P_{ti} *(red crosses) for a target submarine at the nominated depths in the nominated environments with a 'typical' modern sonar system. In each case, accumulation of the cumulative probabilities begins at the maximum range shown on the respective abscissas.*

Figure 6 highlights a well known feature of cumulative probabilities: they never decrease with the passage of time. Some operators find this characteristic disturbing. As an extreme example, the P_d and P_{ti} values attained in Figure 6(d) while passing over the convergence zone beyond 60 km range are maintained through the succeeding long region of very low p_d values. Similarly, P_{ti} in Figure 6(c) plateaus after the bottom bounce opportunity around 18 km although p_d falls to quite low values below 10 km. Clearly, these cumulative probability values will reflect real system performance only if the sonar-system operators have sufficient environmental awareness to know about the convergence-zone and bottom-bounce opportunities. For example, detections will not be made if the operators have selected a range scale that does not extend out as far as the range of the convergence zone or bottom bounce.

As its name indicates, P_{ti} concerns track initiation only; issues of track maintenance and the increase of information staleness with time are not included. Viewed another way, the use of P_{ti} assesses sonar performance in search. Once a track is initiated, the operator may switch to an 'attack' or 'investigate' mode (if available), thereby concentrating acoustic energy in the direction of the target. This would mean, among other things, an alteration in the p_d values.

Another issue concerning the P_{ti} values shown in Figure 6 is the effect of false detections. The calculation resulting in the p_d curves takes a false-detection probability as an input, so false detections have been taken into account at that level. However, in the real world, noise-

induced false detections would also affect track initiation: there is a chance that a false detection could occur sufficiently close to a pair of real detections for it to cause a track initiation that would not otherwise have occurred. The effect is not included in the method described herein, which therefore underestimates P_{ti} by an amount that depends on the false-detection probability. We are undertaking a more elaborate study to explore this [6,7].

3.2.4 Integrating Two or More Sonar Systems

The method of calculating the track-initiation probability for a sonar suite consisting of two or more sonar systems depends on the manner in which the data from the systems are fused. There are several possibilities:

- Each sonar system performs its own tracking and does not pass information to a central location until a track is initiated. This is the simplest case. For systems that act independently of each other in the statistical sense, the track-initiation probability for the suite is given by an expression like Equation (3). That is, for a suite of two sonars,

$$P_{\mathrm{ti,sonar\,suite}} = 1 - \left(1 - P_{\mathrm{ti,sonar1}}\right)\left(1 - P_{\mathrm{ti,sonar\,2}}\right). \quad (5)$$

  The two sonar systems need not, and probably do not, ping at the same time or with the same inter-ping interval. Equation (5) is applied each time either sonar pings to update the $P_{\mathrm{ti,\,sonar\,suite}}$ value.
- The sonar systems send detection information to a centralised tracker. Depending on environmental details, this can result in improved performance over the case of individual tracking, as shown elsewhere [3,11].
- The sonar systems operate fully multistatically. That is, each sonar receives and processes echoes not only from its own pings, but also from those of other sonars in the suite. It is expected that this would result in a further improvement in performance, but the analysis of multistatic systems is complicated and is not considered here.

Figure 7 shows an example of the application of Equation (5). Here, two systems with different characteristics share track-level information. The case shown in Figure 7 is extreme in the sense that neither system by itself achieves a P_{ti} of 0.5 until quite short range, but the suite reaches that value well out. It illustrates the advantage that can accrue even with the simplest type of networking and in the smallest of networks (i.e. a network of just two systems). It also shows how track-initiation probability—whether local or cumulative—provides a way of quantifying networking effects, which single-ping detection probability does not.

3.3 Definition of 'Detection Range'

Use of P_{ti} rather than P_{d} overcomes many of the issues listed on p. 2, but still it is often desired by a client to express the results of the analysis in terms of detection range. When this is the case, there seems to be no avoiding the arbitrary choice of a P_{ti} value (issue 1 on p. 2). Often, a value of 50% is favoured, with inspection of plots like Figures 6 and 7 being used to support the 'reality' of the P_{ti} = 50% rule as a measure of detection range. In our experience,

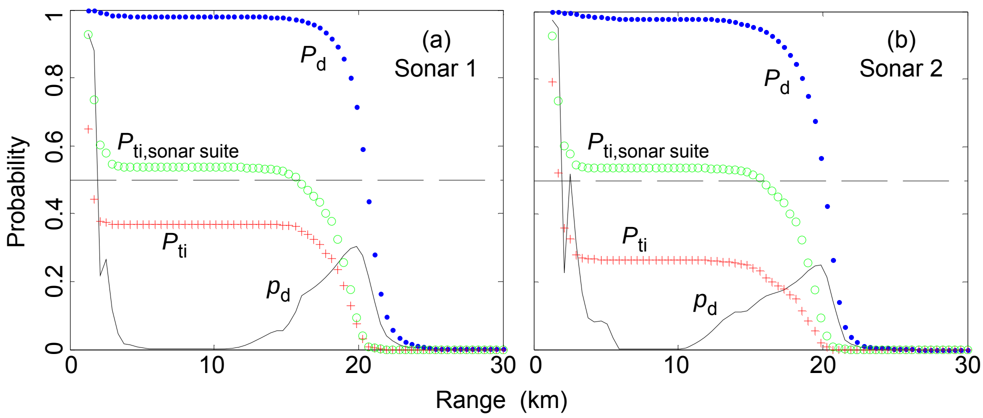

Figure 7: *An example of the effect of track-level fusion of data from two sonar systems in the same environment. The black, red and blue curves have the same meanings as in Figure 6, the green curve (which is the same in both panels) gives values of* $P_{ti,\ sonar\ suite}$. *The broken line shows the 50% probability value, to indicate the range at which* $P_{ti,\ sonar\ suite} = 0.5$.

operators and other clients seem to feel more comfortable with the 50% level than any other. It is as though 50% is regarded as more natural than other values. Is it, perhaps, that the use of a value other than 50% draws unwelcome attention to the arbitrariness of the whole process? This touches on the psychology of decision making, which is for other studies to address.

Figure 6 also provides practical support for a detection-range definition using the 50% value in that, for these four cases, the range at which P_{ti} = 50% is little different from that at which P_{ti} equals, say, 90%. This is because the P_{ti} curve rises so sharply in the relevant region. That is, the resulting detection range is insensitive to the P_{ti} value chosen, within reasonable limits. This feature is, however, not true for all combinations of sonar system, environment, time of year, target type and depth. As an extreme example, Figure 7 shows a case where the value of $P_{ti,\ sonar\ suite}$ hovers only a little above 50% for an extended range. It is also a case where forming a network results in a much larger detection range than either sonar system is capable of achieving on its own.

It may be questioned whether a range based on P_{ti} should be termed a 'detection range'. Would not 'track-initiation range' be more appropriate? Possibly, but the client may be unwilling to accept an unfamiliar term. We can rationalise the use of 'detection range' in terms of the difference noted in §2.1 between engineering and operator interpretations of 'detection'. A strict engineering point of view would indeed hold that the detection range should be defined in terms of P_d rather than P_{ti}. However, operators understand 'detection' to include aspects of identification and the filtering out of false detections. Application of a track-initiation rule involves elements of the second, so providing a level of justification for the usage.

4. Summary and Conclusion

In this note, we propose the use of cumulative track-initiation probability P_{ti}, rather than single-ping or cumulative detection probability, as a basis for metrics for assessing active-sonar-system performance in anti-submarine warfare. The definition of P_{ti} is: the probability that a track is initiated on the current ping, or has already been initiated on a previous ping.

The issues, advantages and disadvantages of P_{ti} compared with the other probabilities (including local p_{ti}) are summarised schematically in Table 1. In detail:

- Cumulative probabilities, whether P_{ti} or P_{d}, require a scenario for their calculation. Hence their values are scenario dependent to a greater extent than single-ping or local probabilities. This can be mitigated somewhat by concentrating on scenarios simple enough to form basic tactical building blocks of sufficient generality to be useful to a principal warfare officer faced with the task of formulating a detailed tactical response to the contingency at hand.
- Cumulative probabilities provide a basis for comparing sonar systems with different ping rates or transmission modes, unlike single-ping or local probabilities.
- Single-ping p_{d}, by definition, does not allow the analysis of the effect of a long series of pings, unlike the cumulative probabilities. Local p_{ti} considers groups of 5 pings only (or n pings where the track-initiation rule is m-in-n).
- Track-initiation probabilities provide some recognition of the problem of false detections, unlike detection probabilities, while skirting around the difficult issue of data association.
- As regards several sonar systems acting together and pooling data, the single-ping detection probability p_{d} provides no basis for analysis. Combinations of several p_{d} values are required, as is the situation for the other three probabilities. However, P_{d} permits the analysis of track-level fusion only, whereas p_{ti} and P_{ti} allow in addition the analysis of detection-level data fusion.
- Since cumulative probabilities never decrease with time, they provide an unambiguous basis for defining detection range, unlike local probabilities.
- The actual definition of detection range requires the choice of a probability value, which is unavoidably arbitrary. All probabilities have this deficiency, but it is worst for cumulative detection probability P_{d}, because the 'natural' choice of 50% — the one value that an operator is unlikely to query — gives unrealistically large detection ranges in most environments. On the other hand, experience suggests that P_{ti} = 50% works well, at least for persistent sonar systems like ship-borne hull-mounted sonars and towed arrays.

Finally, we note that the purpose of any piece of military operations analysis is to inform a decision maker. The success of the analysis is closely related to how useful the client finds it. Analytical techniques based on cumulative track-initiation probability have been used to inform decisions in support of a major defence acquisition project, to the satisfaction of the client.

Table 1: Summary of strengths and weaknesses of various probabilities for the purposes of defining the 'detection range' of a sonar system

Probability:	single-ping detection probability	local track-initiation probability	cumulative detection probability	cumulative track-initiation probability
independent of scenario*				
allows comparison of sonars with different ping rates or modes				
provides some recognition of the influence of false detections				
includes effect of many repeated pings		5 pings only		
provides a basis for analysing networks of sonar systems			track-level fusion only	
unambiguous definition of detection range				
no arbitrariness in definition of detection range	p_{d} = 50% acceptable	p_{ti} = 50% acceptable	P_{d} = 50% unrealistic	P_{ti} = 50% acceptable

*This refers to the disposition and motion of the platforms in the scenario. All probabilities are absolutely dependent on the environment, time of year, nature of the target and target depth.

Acknowledgement and Disclaimer

We thank Paul Clarke for commenting on a draft of this note and for carrying out the sonar-performance calculations that provided the illustrative p_{d} values (shown in Figures 1, 3, 6 and 7) on which the calculation of the other probabilities is based. Paul has sufficient experience in modelling the sonar performance of real systems to be able to ensure that the p_{d} values used in this work display realistic features, but do not represent the actual performance of any known sonar system either currently in service or proposed for the foreseeable future.

References

[1] P.H. Dahl, J.H. Miller, D.H. Cato & R.K. Andrew (2007) 'Underwater ambient noise', *Acoustics Today*, vol 3, pp. 23–33.

[2] F.G. Phillips, R.S.W. Allen & M. Spillane (1998) 'Analysis of anti-submarine warfare systems', Technical report DSTO-TR-0713 of the Defence Science and Technology Organisation (CLASSIFIED).

[3] M.P. Fewell, J.M. Thredgold & D.J. Kershaw (2008) 'Benefits of sharing detections for networked track initiation in anti-submarine warfare', Technical report DSTO-TR-2086 of the Defence Science and Technology Organisation.

[4] A. Davies & M. Spillane (1997) 'A method for optimising sonar ping rates', Annex L in A. Davies & M. McKenzie 'Anti-submarine warfare capability study', document CAB ASWCS 2/97 of the Capability Analysis Branch, Strategic Policy & Plans Division, Australian Defence Headquarters.

[5] A. Khan & S. Bocquet (2009) 'A maritime patrol model for comparison of AIR 7000 candidate aircraft', Technical note DSTO-TN-0901 of the Defence Science and Technology Organisation (CLASSIFIED).

[6] J.M. Thredgold, S.J. Lourey, H.X. Vu & M. P. Fewell (2009) 'A simulation model of networked tracking for anti-submarine warfare', Technical report DSTO-TR-2372 of the Defence Science and Technology Organisation.

[7] J.M. Thredgold & M. P. Fewell (2009) 'Distributed versus centralised tracking in networked anti-submarine warfare', Technical report DSTO-TR-2373 of the Defence Science and Technology Organisation.

[8] S. Coraluppi & C. Carthel (2005) 'Distributed tracking in multistatic sonar', *IEEE Transactions on Aerospace and Electronic Systems*, vol. 41, pp. 1138–47.

[9] F.R. Castella (1976) 'Sliding window detection probabilities', *IEEE Transactions on Aerospace and Electronic Systems*, vol. AES-12, pp. 815–819.

[10] P. Williams (1998) 'Evaluating the state probabilities of M out of N sliding window detectors', Technical note DSTO-TN-0132 of the Defence Science and Technology Organisation.

[11] M.P. Fewell, J.M. Thredgold & D.J. Kershaw (2008) 'Sharing detections among networked sonars—a practical way forward for anti-submarine warfare?', *Journal of Battlefield Technology*, vol. 11, no. 2, pp. 17–22.

DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION DOCUMENT CONTROL DATA					1. PRIVACY MARKING/CAVEAT (OF DOCUMENT)
2. TITLE Cumulative Track-Initiation Probability as a Basis for Assessing Sonar-System Performance in Anti-Submarine Warfare			3. SECURITY CLASSIFICATION (FOR UNCLASSIFIED REPORTS THAT ARE LIMITED RELEASE USE (L) NEXT TO DOCUMENT CLASSIFICATION) Document (U) Title (U) Abstract (U)		
4. AUTHOR(S) M.P. Fewell and J.M. Thredgold			5. CORPORATE AUTHOR DSTO Defence Science and Technology Organisation PO Box 1500 Edinburgh South Australia 5111 Australia		
6a. DSTO NUMBER DSTO-TN-0932	6b. AR NUMBER AR-014-676		6c. TYPE OF REPORT Technical Note	7. DOCUMENT DATE December 2009	
8. FILE NUMBER 2009/1083266	9. TASK NUMBER NAV 07/096	10. TASK SPONSOR COMFLOT	11. NO. OF PAGES 15	12. NO. OF REFERENCES 11	
13. URL on the World Wide Web http://www.dsto.defence.gov.au/corporate/reports/DSTO-TN-0932.pdf			14. RELEASE AUTHORITY Chief, Maritime Operations Division		
15. SECONDARY RELEASE STATEMENT OF THIS DOCUMENT *Approved for public release* OVERSEAS ENQUIRIES OUTSIDE STATED LIMITATIONS SHOULD BE REFERRED THROUGH DOCUMENT EXCHANGE, PO BOX 1500, EDINBURGH, SA 5111					
16. DELIBERATE ANNOUNCEMENT No Limitations					
17. CITATION IN OTHER DOCUMENTS Yes					
18. DSTO RESEARCH LIBRARY THESAURUS http://web-vic.dsto.defence.gov.au/workareas/library/resources/dsto_thesaurus.shtml Sonar performance, Submarine detection, Anti-submarine warfare, Detection range, Track-initiation probability					

19. ABSTRACT

Calculations of sonar-system performance for anti-submarine warfare yield detection probability as a function of range to the target; whereas sonar operators typically think in terms of a 'detection range'. This Note considers how to connect these two concepts for active sonar. Four types of probability are explored and their relative advantages teased out in the light of real-world examples. It is concluded that cumulative probability of track initiation provides the most practical route to a definition of detection range, for two reasons. First, unlike detection probability as usually computed, it produces values that operators are likely to regard as believable. Secondly, it reflects more closely than the standard detection probability the steps that operators typically go through in deciding to declare a detection.